REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
May 1998

3. REPORT TYPE AND DATES COVERED
Professional Paper

4. TITLE AND SUBTITLE
SHPL: The Scaleable High-Performance LAN on the CVN 71

5. FUNDING NUMBERS
PR: CS23
PE: 0603513N
AN: DN288625

6. AUTHOR(S)
L. M. Blackwell, P. D. Adams, C. Hansen, and R. A. Fastring

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Space and Naval Warfare Systems Center
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Naval Sea Systems Command
Washington, DC 20362

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

19990630 084

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENTR
Approved for public release; distribution is unlimited.

13. ABSTRACT *(Maximum 200 words)*

This paper reports on the Scaleable High-Performance LAN (SHPL), the networking infrastructure for the Navy's FY-96 Joint Power Projection/Real Time Support (JPP/RTS) Core Technology Program aboard USS Theodore Roosevelt (CVN 71). The program demonstrates advanced concepts in strike planning, visualization, and execution using the latest technology in high-speed computing, 3-D graphics, and networking. The Joint Power Projection application that SHPL supports involves the coordination of Air Tasking Orders (ATOs) among CTAPS, TSCM and APPEX, the display via APPEX of TSCM auto-route missions, the review of missions by Strike Lead, the passing of missions to TAMPS for detailed aircraft planning, the review of the mission concept by CAG, and the overall mission preview by TSCM and APPEX. The SHPL, or networking element of the JPP/RTS Core Technology Program, is a totally fiber optic shipboard network that is based on asynchronous transfer mode (ATM) and switched-Ethernet technology. The SHPL network demonstrates both permanent virtual circuits (PVCs) and switched virtual circuits (SVCs), handles IP frames via both classical IP (RFC 1577) and FORE IP, and demonstrates the LAN emulation functionality of the ATM Forum's LANE 1.0 standard..

In addition to handling computer data such as file transfer, imagery and email, the SHPL network handles voice and video in the form of interactive video teleconferencing (VTC), shared X-windows applications via SharedX, interactive electronic whiteboarding, and NTSC video via ATM. SHPL also demonstrates the network management of a hybrid ATM and Switched-Ethernet network, and cable plant configuration management. The use of SHPL on the CVN 71 demonstrates what high-performance networking technology can do for next generation Navy ships, and how networking can extend the useful life of current ships.

Published in Proceedings of the 11th Ship Control System Symposium, April 1997.

14. SUBJECT TERMS
Mission Area: Communications
Fiber Optic Scaleable High-Performance LAN (SHPL)
Joint Power Projection/Real Time Support (JPP/RTS) Core Technology Program
Networking

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
			Same as Report

NSN 7540-01-280-5500

Standard form 298 (FRONT)

21a. NAME OF RESPONSIBLE INDIVIDUAL	21b. TELEPHONE *(Include Area Code)*	21c. OFFICE SYMBOL
P. D. Adams	(619) 553-3422	D8205

SHPL: The Scaleable High-Performance LAN on the CVN 71

L. M. Blackwell, P. D. Adams, C. Hansen & R. A. Fastring
(Blackwell) Naval Sea Systems Command, Code 03J2, National Center Building 4, Room 428, 2531 Jefferson Davis Highway, Arlington, VA 22242, Blackwell_Luther_M@hq.navsea.navy.mil
(Adams and Hansen) Naval Command, Control and Ocean Surveillance Center, RDT&E Division, Code 8205, 271 Catalina Boulevard, San Diego, CA 92152, adamspd@nosc.mil
(Fastring) Synetics Corporation, Suite 210, 9449 Balboa Avenue, San Diego, CA 92123, richard_fastring@synetics.com

Abstract

The Scaleable High-Performance LAN (SHPL) is the networking infrastructure for the Navy's FY-96 Joint Power Projection / Real Time Support (JPP/RTS) Core Technology Program aboard *USS Theodore Roosevelt* (CVN 71). The program demonstrates advanced concepts in strike planning, visualization and execution, using the latest technology in high-speed computing, 3-D graphics, and networking.

The Joint Power Projection application that SHPL supports involves the coordination of Air Tasking Orders (ATOs) among CTAPS, TSCM and APPEX, the display via APPEX of TSCM auto-route missions, the review of missions by Strike Lead, the passing of missions to TAMPS for detailed aircraft planning, the review of the mission concept by CAG, and the overall mission preview by TSCM and APPEX. Individual aircraft tracks can be pulled at any time by TOPSCENE for pilot and weapon systems officer rehearsal.

The SHPL, or networking element of the JPP/RTS Core Technology Program, is a totally fiber optic shipboard network that is based on asynchronous transfer mode (ATM) and switched-Ethernet technology. SHPL supports high-speed (155Mbps) ATM data transfer among Sun workstations, Silicon Graphics workstations, and Hewlett Packard (TAC-3/4) workstations, and also supports ATM to legacy LAN data transfer between these workstations and fiber optic Ethernet workstations. The ATM-to-ATM data transfer uses FORE Systems ASX-200BX ATM switches, and the ATM-to-LAN conversion takes place in Cisco Systems Catalyst 5000 edge devices.

The SHPL network demonstrates both permanent virtual circuits (PVCs) and switched virtual circuits (SVCs), handles IP frames via both classical IP (RFC 1577) and FORE IP, and demonstrates the LAN emulation functionality of the ATM Forum's LANE 1.0 standard.

In addition to handling computer data such as file transfer, imagery and email, the SHPL network handles voice and video in the form of interactive video teleconferencing

(VTC), shared X-windows applications via SharedX, interactive electronic whiteboarding, and NTSC video via ATM.

SHPL also demonstrates the network management of a hybrid ATM and Switched-Ethernet network, and cable plant configuration management. This paper summarizes the SHPL architecture on the *USS Theodore Roosevelt* (CVN 71), presents lessons learned, and provides a view of the direction that shipboard networking might take in the future.

The use of SHPL as the networking core of Joint Power Projection on the CVN 71 demonstrates what high-performance networking technology can do for next generation Navy ships, and how networking can extend the useful life of current ships. While the immediate application was strike planning, SHPL could just as easily be applied to ship control functions such as steering, damage control, machinery control, etc.

1 Introduction

This paper summarizes the networking aspects of the Joint Power Projection / Real Time Display (JPP/RTS) Core Technology Program demonstration that took place on the aircraft carrier *USS Theodore Roosevelt* (CVN 71) during the period of July through September 1996. The network that supported JPP/RTS was an OC-3c (155Mbps) asynchronous transfer mode (ATM) network with switched-Ethernet edge devices. This network, called the Scaleable High Performance LAN (SHPL), provided a backbone network for the interconnection of all workstations that have a direct, or indirect, interface with the strike planning function. These include high performance workstations that process the maps and high resolution imagery involved in strike planning, and also the personal computers that are used for administrative functions such as email.

2 The SHPL Architecture

A top level diagram of the SHPL architecture is shown in Figure 1. At the core of SHPL are six ATM (asynchronous transfer mode) switches AS1, AS2, ..., AS6. Each is a Fore Systems ASX200BX ATM switch that is configured with four 4-port OC-3c interface cards so as to provide sixteen 155Mbps connections to the switch. On each switch, 4 of the 16 ports are used for switch-to-switch trunklines and 12 are used for connecting external devices to the switch. On two switches one of the OC-3c ports is used to connect a Cisco Catalyst 5000 edge device which is used to interconnect between the ATM network and legacy Ethernet local area network (LAN) devices. The remainder of the 155Mbps ports are used for the connection of high performance workstations which have OC-3c network interface controller (NIC) cards installed.

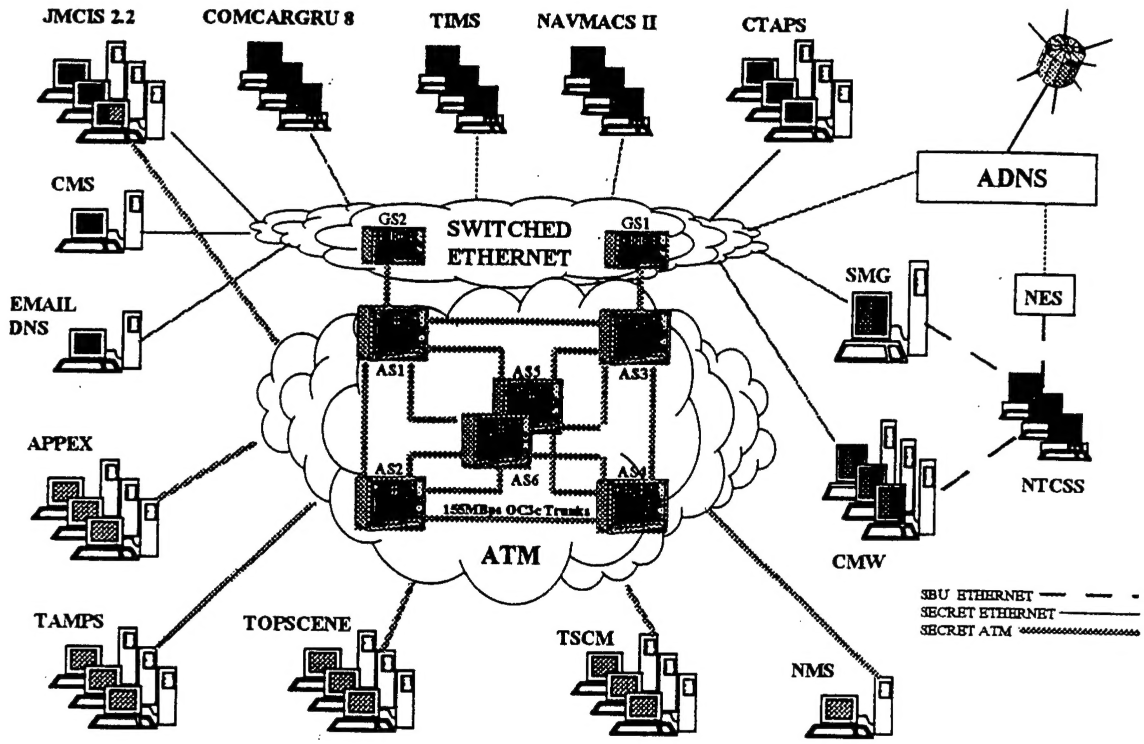

Figure 1. The SHPL Architecture on USS Theodore Roosevelt

Abbreviations used in the SHPL diagram:

ADNS - Advanced Digital Network System
APPEX - Advanced Power Projection Planning and Execution System
AS1, ... ,AS6 - FORE Systems ASX-200BX 16-port ATM switches
CMS - Cable Configuration Management System
CMW - Compartmented Mode Workstation
COMCARGRU - Commander Carrier Group
CTAPS - Contingency Theatre Air Control Automated Planning System
DNS - Domain Name Service
GS1,GS2 - Cisco Catalyst 5000 ATM edge devices
JMCIS - Joint Maritime Command Information System
NAVMACS - Naval Modular Automated Communications System
NES - Network Encryption System
NMS - Network Management System
NTCSS - Navy Tactical Command Support System
SMG - Standard Mail Guard
TAMPS - Tactical Aircraft Mission Planning System
TIMS - Tactical Information Management System
TOPSCENE - Tactical Operational Scene
TSCM - Tactical Strike Coordination Module

The workstations connected to SHPL via 155Mbps OC-3c links include Hewlett-Packard TAC-4 workstations using HPA-200E/OC3ST EISA-bus NICs, SGI Onyx workstations using VMA-200/OC3ST-SGI VME-bus NICs, SGI Impact and Indigo-2 workstations using ESA-200E/OC3ST EISA-bus NICs, SGI Indy workstations using GIA-200E/OC3ST-S GIO-bus NICs, Sun workstations using SBA-200E/OC3ST Sbus NICs, and a Windows NT Intel-based workstation with a PCA-200PC/OC3ST PCI-bus NIC. All of these 155Mbps NICs were Fore Systems cards with on-board 25mhz i960 RISC processors.

These 155Mbps workstations supported functions such as APPEX, TAMPS, TOPSCENE, TSCM, and NMS. (Refer to the table below Figure 1 for the definition of these acronyms.)

In addition to these workstations that connected to the network at the OC-3c level, a number of workstations connected at the 10Mbps Ethernet level, either because they had not yet transitioned to a higher performance interface, or had no need for higher speeds. The Ethernet workstations served functions such as JMCIS, COMCARGRU 8, TIMS,

NAVMACS II, CTAPS, SMG, CMW, CMS and EMAIL/DMS. (Refer to Figure 1 for definitions.) These Ethernet workstations interfaced to SHPL via one of the two Cisco Catalyst 5000 edge devices. All of the Ethernet ports on the Catalyst 5000 devices were fiber (10Base-FL) switched-Ethernet ports. Some of these switched-Ethernet ports were dedicated to a single Ethernet workstation (EMAIL/DNS, CMS, NAVMACS II, SMG, and CMW) while in other cases an existing shared-bandwidth Ethernet LAN was connected to the switched-Ethernet port (JMCIS, COMCARGRU 8, TIMS, and CTAPS).

2.1 MultiLevel Secure (MLS) Aspects of SHPL

SHPL itself is a system high GENSER SECRET network, but it is connected to a sensitive but unclassified (SBU) NTCSS network via a Standard Mail Guard (SMG) and a number of Compartmented Mode Workstations (CMWs). The SMG enables unclassified email (without attachments) to be exchanged between the two different security levels. Since the NTCSS network is connected via ADNS to satellite facilities, this SMG linkage enables personnel at GENSER SECRET workstations to send/receive unclassified email to/from worldwide correspondents via the Internet/NIPRNET.

The compartmented mode workstations (CMWs) are Hewlett Packard TAC-4 workstations that run a secure MLS version of the UNIX operating system. This MLS operating system enables an operator to open multiple windows on his workstation, with each window at a different security level. In one window he can be sending/receiving unclassified email via the Internet, while in another window viewing a secret JMCIS tactical picture or browsing the secret SIPRNET.

2.2 The SHPL Protocol Configuration

With the exception of the TV distribution using the AVA-300 (see below), all applications running on the SHPL network used the IP protocol. Two separate IP subnets were configured: One was used for LAN Emulation using the ATM Forum's LANE 1.0 standard, and the other was used for Fore Systems FORE IP/SPANS (Simple Protocol for ATM Network Services) protocol.

FORE IP/SPANS is functionally similar to RFC-1577 (Classical IP and ARP over ATM) except that it supports IP multicasting, provides a more robust network-to-network (NNI) routing/signaling interface, and supports direct communication of all hosts on a physical ATM network without the use of IP routers.

Switched Virtual Circuits (SVCs) using UNI 3.0 signaling were used for all connections with the exception of the use of Permanent Virtual Circuits (PVCs) for connections to the AVA-300 video-to-ATM converter. ATM Adaptation Layer 5 (AAL-5) was used for all connections.

For video distribution, motion JPEG (M-JPEG) compression was used both within RTDS for TV distribution and within Insoft's Communiqué! for videoteleconferencing. Using Communiqué! VTC at a frame rate of 30 frames per second and a QIF (quarter size) frame resulted in a data rate between 700 kbps and 1 Mbps for each direction of video transmission. (By way of comparison, broadcast quality video using M-JPEG at a Q-factor of 20 for the transmission of full size, full rate NTSC video results in a data rate of approximately 18 Mbps.) For audio conferencing, VAT was used with the PCM option, e.g., 78 kbps 8-bit mu-law encoding at 8k Hz sampling rate with 20 ms frames.

2.3 Reliability and Survivability Considerations

The ATM switches and the Ethernet-ATM edge devices all have dual, hot-swappable power supplies with separate line cords. One line cord is connected to a TrippLite line conditioner and the other line cord is connected to an American Power Conversion UPS. The UPS and line conditioners protect against ground faults, over-voltage, brownouts and spikes. Each of the four SHPL racks is fed from a different circuit via a different breaker. The ATM switches and edge devices have hot-swappable line cards. Each ATM switch has four trunklines to four different switches. All SHPL racks are located in air-conditioned spaces and are isolated from shock/vibration by coil springs beneath the racks and between the top of each rack and the nearest bulkhead.

3 SHPL Measured Performance

3.1 Throughput

For ATM-to-ATM transfers of large (60 MB) files via ftp over LAN emulation, the throughput varied from 5.6 Mbps to 40.8 Mbps as a function of the driver/workstation pairing. Note that this does not represent the maximum throughput that can be sustained by a 155 Mbps OC-3c link through the network inasmuch as it includes the limitations of the workstation's operating system, the ftp protocol, and the NIC driver software.

3.2 Broadcast Rate

The broadcast rate on the SHPL backbone averaged 19.2 packets/sec with a peak of 75 packets/sec.

3.3 Cell Rate

The cell rate on the most heavily loaded ATM switch typically peaked at less than 5000 cells/sec, which is a very small fraction of the peak throughput per switch of 2.5 Gbps.

3.4 Dropped Cells

Over a 24 hour period during which up to 42 million cells were transferred, the ForeView network management program reported zero dropped cells.

3.5 Maximum Potential Network Bandwidth

Assuming that a pair of high-performance workstations with high-performance NICs could achieve a throughput of 100 Mbps in each direction over a full-duplex OC-3c ATM circuit, then the network bandwidth of SHPL on *USS Theodore Roosevelt* is 4.8 Gbps for all non-local connections (as limited by the 24 inter-switch trunklines), or 14 Gbps for all local connections (as limited by the 70 ports to user devices exclusive of the two Catalyst 5000 edge devices). Typically, some connections will be local (both users connected to the same switch) and some connections will be non-local (each user connected to a different switch) so that the effective SHPL bandwidth is between 4.8 Gbps and 14 Gbps.

3.6 Backplane Bandwidth

The bandwidth of the backplane in the Fore Systems ATM switches is 2.56 Gbps. The Cisco Catalyst 5000 edge devices have a backplane bandwidth of 1.2 Gbps.

3.7 UPS Run Time following Failure of AC Power

Of the four racks of SHPL equipment, two could survive on battery power from the UPS for 71 minutes and two could survive for 21 minutes (the difference being based on the amount of equipment in the different racks).

4 SHPL Multimedia Applications

The primary purpose of SHPL is to provide a high-bandwidth interconnection for Navy tactical applications such as JMCIS, TAMPS, APPEX, TSCM, and TOPSCENE. However, the available bandwidth exceeds the immediate needs of these tactical applications, and this extra bandwidth can be used to improve the user-to-user collaboration and the outside-world awareness of the user via COTS (commercial off-the-shelf) multimedia applications.

It should be noted that multimedia applications are generally more mature in the personal computer arena (running on Windows, Windows for Workgroups and Windows 95) than in the tactical (UNIX) computer arena. However, the tactical applications that ran on SHPL were hosted on TAC-4, Silicon Graphics and Sun workstations, so all of the multimedia applications that were demonstrated on *USS Theodore Roosevelt* were UNIX-based applications.

Some of these multimedia applications are shown in Figure 2. The left side of this figure shows how NTSC television signals can be digitized, compressed (via Motion JPEG), and converted to ATM cells by an AVA-300 video-to-ATM converter (manufactured by Nemesys Research Ltd. and OEMed by Fore Systems), and distributed throughout the network. These TV signals can come from any television signal source, such as off-the-air broadcasts (e.g., CNN News), remote feeds from unmanned aerial vehicles (UAVs), flight deck cameras that monitor aircraft launching and recovery, security cameras or training videos. This digitized video can be converted back to analog TV signal format in an ATV-300 ATM-to-video converter for display on a TV monitor such as the large screen display in *Theodore Roosevelt's* War Room. Alternatively, the digitized video can be fed via the ATM backbone into a workstation (e.g., TAC-4) that is equipped with a video board (e.g., Parallax Graphics) and which runs Nemesys Research's Real Time Display Software.

On the right side of Figure 2, one can see that the ATM network can be used for video-teleconferencing (VTC) and application sharing via electronic white boards or X-windows sharing..

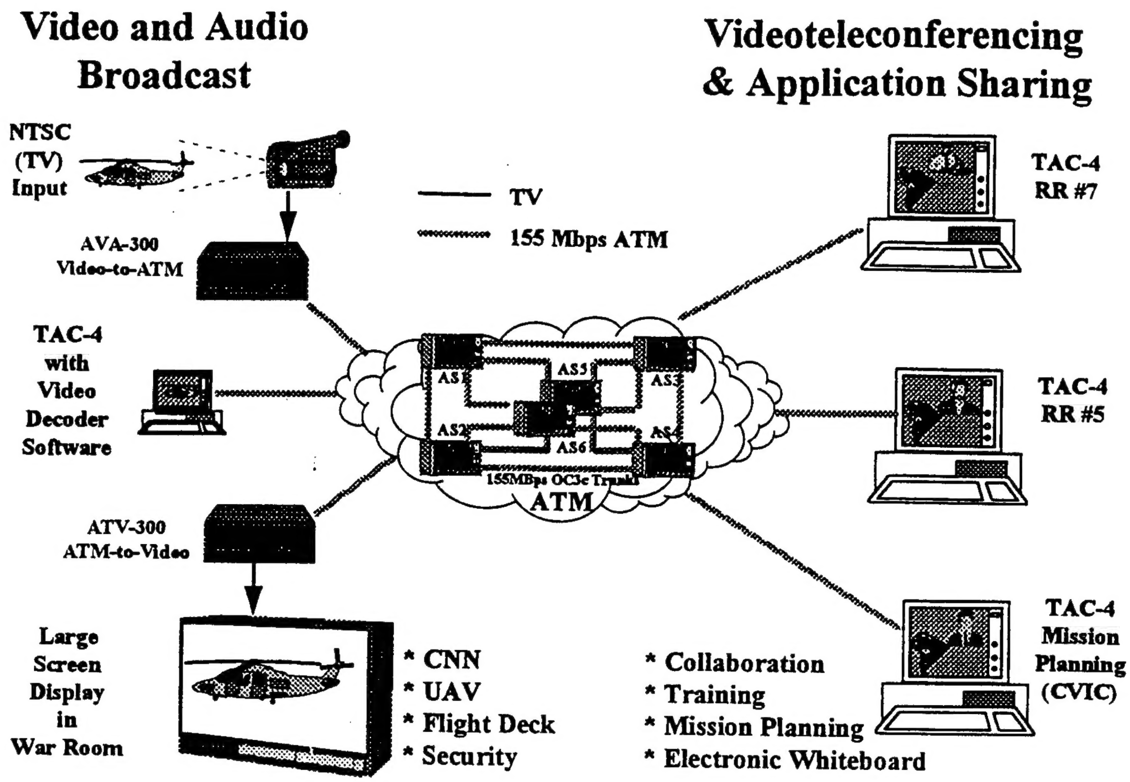

Figure 2. SHPL Multimedia Applications

A typical screen layout for a multimedia terminal is shown in Figure 3. In the upper left is CNN News displayed in a Real Time Display Software (RTDS) window. The upper center has an X-Windows application that is running on one workstation and is being displayed on another workstation via SharedX (Hewlett Packard's shared X-window utility for the HP-UX operating system). The shared application can be some administrative tool, such as a calendar program, or it can be a tactical application such as JMCIS or APPEX. The application will normally be running on a UNIX workstation, but it can be shared with either another UNIX workstation or with a PC workstation that is running X-Windows. This latter option enables personnel who prefer a PC workstation as their desktop productivity tool to also stay "plugged-in" to the tactical situation by displaying, and optionally controlling, a tactical application via a SharedX screen.

The upper right shows a videoteleconference (VTC) that is shown via Insoft's Communiqué!. The lower right has an X-Host Chat (XHCHAT from the University of Karlsruhe) real time text session between two operators at different workstations. The lower left corner shows the audio control panel from Visual Audio Tool (VAT from UC Berkeley's Lawrence Berkeley Laboratory) that is being used for two-way voice-over IP. In the center of the screen two operators are interacting via Insoft's Communiqué! White Board, which is an electronic white board upon which both operators can draw, type, or annotate an image file that has been captured from some other application.

The VAT, XHCHAT, and Communiqué! applications as implemented on *USS Theodore Roosevelt* were two-party connections due to protocol limitations of the HP-UX 9.0 operating system running on the TAC-4 machines. However, with multicast IP (available under HP-UX 10.0), these could be multiparty connections with three or more participants.

(The SharedX application does not require multicast IP and can be used to share an application among three or more workstations using the 9.0 version of HP-UX.)

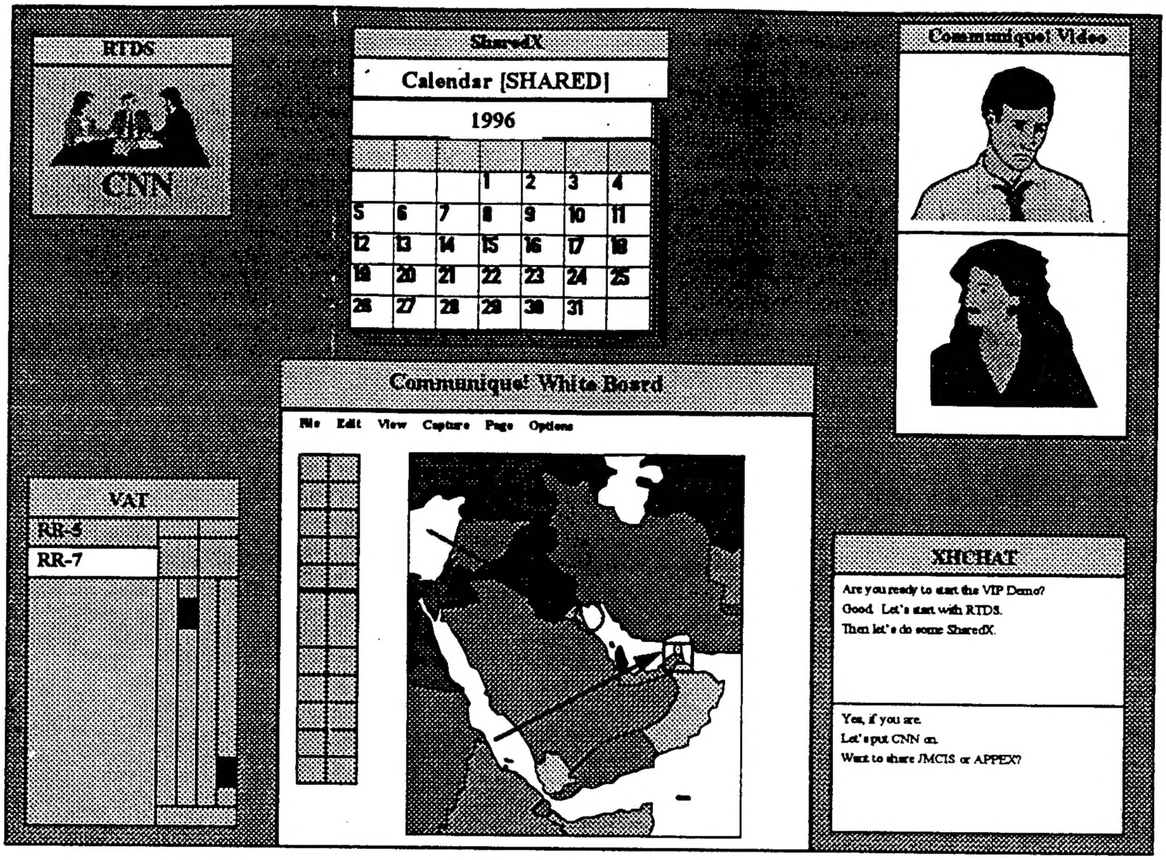

Figure 3. A Typical Multimedia Screen

5 Lessons Learned

5.1 Integration of Multiple Software Applications on a Common Workstation

All Navy ships have space problems. As more systems are brought aboard, they compete for ever decreasing space. Therefore the integration of multiple applications into common workstations is highly desirable. However, existing Navy software applications that were written for standalone, dedicated workstations cannot necessarily be ported to a common workstation without sometimes significant software development.

Furthermore, each development manager has a responsibility to maintain configuration control over his application, not only for issues such as which directories files are placed in, what permissions are setup for reading, writing or executing files, etc., but also sometimes for the "look and feel" of the application in terms of color selections, the size and placement of windows, which windows are prioritized to be always "on top", etc. If multiple applications are integrated into a common workstation, there is no guarantee that each successive software integrator will, or even could, maintain the configuration control desired by earlier development managers.

This problem of multiple applications sharing the same workstations was observed during the demonstration of SHPL applications on *USS Theodore Roosevelt*. A demonstration of SHPL multimedia applications (distribution of digital TV, videoteleconferencing, voice-over-IP, shared electronic white boards, shared X-windows, etc.) was needed, but there was

insufficient space to add new workstations to the ship to demonstrate these applications. Instead, two workstations that were configured for another tactical application were made available for the multimedia demonstration. However, the root password on these machines was not available during the pre-demonstration setup, which lead to only partial success in configuring the multimedia software to run on these machines. Furthermore, some instabilities noted during the demonstration were (correctly or incorrectly) attributed to the inability to configure the workstations exactly as desired. Thus, one "lesson learned" is *never use another project's workstations if you do not have complete configuration control.* However, this is the exact opposite conclusion to that which is desired: application integration using common workstations.

The above problem is serious enough at the "tactical application" level where some degree of design control and configuration management is effective. But many, if not all, users desire to perform "productivity" tasks such as wordprocessing, viewgraph presentations, Internet browsing, spreadsheet development, etc. on this same machine if it is to be his common workstation. (This desire was previously not achievable when tactical computers and personal computers were radically different. However, Microsoft Windows can now be emulated on tactical UNIX machines, and the Navy is slowly moving toward the use of Windows 95 or Windows NT for tactical applications.) In other words, the line between personal computer and tactical display is blurring. If there is no space for multiple machines, then the common workstation not only has to integrate multiple tactical applications, but perform the personal computer role as well. And at that point all attempts to rigidly control the software load configuration on the common workstation falls apart. Few, if any, professionals will accept a "standard configuration" personal computer. Everyone needs, or at least really believes that he needs, special software, or software configured in a special way, to get his job done. So it appears that flexibility in software configuration for common workstations will have to be accommodated, and new standards for the minimum configuration controls that are truly needed for interoperability must be defined without stifling the user's desire to configure, and continuously reconfigure, his workstation.

5.2 Configuring VAT

When the Visual Audio Tool (VAT) was configured properly, the speech quality was excellent. However, there are a large number of options, plus multiple slider controls, all of which significantly impact voice quality and some of which are rather critical in their adjustment. It was also noted that, for as yet unexplained reasons, a two-way VAT session would randomly lock into a mode in which the audio continued to be transferred but the slider controls could not be moved.

Some, but not all, of the desired configuration parameters could be preset by appropriate command line options when the VAT program is started. However, since some of the adjustments must be made after the program starts up, it is doubtful if Navy personnel will make extensive use of the current version of VAT in true tactical situations in which there is insufficient time (or motivation) to fine-tune the parameters.

5.3 Using Insoft's Communiqué!

The video and shared white board functions of Communiqué! worked very well during this demonstration. However, the voice quality provided by Communiqué! was poor and VAT was used as a better alternative. (The problem with Communiqué!'s voice quality was

traced to their reliance on the audio server (Aserver) function within the TAC-4 which did not appear to be stable in the software version that was used.)

5.4 Two-way versus Multiparty Conferences

Because the TAC-4 workstations were running HP-UX 9.0, which has no support for IP multicast, the demonstration was limited to two-way conferences with VAT, VTC and electronic white board. These functions would be much more useful if they could be extended to three or more participants. This should be possible with HP-UX 10.0 with IP multicast.

5.5 Using Hewlett Packard's SharedX

From a practical standpoint, SharedX worked with some, but not all, applications. The electronic white board function within Communiqué! worked extremely well at a remote workstation acting as an X-windows server. However, when the JMCIS application was shared to this same workstation, it worked only partially. The basic map would display at the remote workstation, but the remote workstation could not activate any of the options on the pull-down menus. (It is suspected, but not known, that the JMCIS software spawns new UNIX processes for each pull-down menu selection, and these new processes are not shared automatically when the top-level screen is shared.)

We also noted that the APPEX Time Line, when shared via SharedX, slowed down the cursor movement on both the host machine running the APPEX code and the remote X-windows server machine so as to make the application useless. The lesson to be learned here is that SharedX applications have to be written with X-windows sharing in mind. Many applications that were written to run on a single standalone UNIX machine may not operate properly from a remote X-window server.

5.6 Reliability of COTS Networking Hardware aboard Navy Ships

This demonstration clearly showed that when high-quality commercial off-the-shelf (COTS) networking components are properly installed (mounted in shock/vibration isolated ruggedized racks, with rack-level blowers, power line conditioners and/or uninterruptable power supplies (UPS), proper cable ties and strain relief, etc.), the result can be a reliable full-time network. There were no hardware failures of the SHPL network during the entire test period. However, see below:

5.7 Reliability of COTS Networking Software aboard Navy Ships

Some of the user workstations experienced brief periods of partial network failure (inability to gain access to other workstations via LAN emulation) that were subsequently rectified by a software upgrade in the ATM switches and network interface cards. This situation has been noted before. Software-based controllers and flash PROMS enable network vendors to rush new technology to market before the controlling software is mature, and the system integrator can count on one or two software upgrades before the system works as advertised. This is something to be accounted for in program schedules. Given the option of using obsolete hardware with unknown future support, or bleeding edge technology with immature software, the Navy is better served by the latter since the latter option eventually stabilizes for a period before becoming obsolete.

5.8 Keeping Up-to-date with Software Versions

One of the tactical applications had been written to run on Sun workstations using Solaris 2.3. The ATM network interface controller cards (NICs) had drivers that worked only with Solaris 2.4 and beyond. The NIC vendor had little incentive to backwards-port the drivers to an older operating system which would provide poorer performance and which was little used in the commercial world. The tactical application had to be ported (in a rush) to the newer operating system in order to interface via ATM. The lesson learned here is to keep up with the commercial world if you intend to use COTS. You cannot afford to be one or two revisions back in your compiler, operating system, network drivers, etc. and expect interoperability.

6 Improvements for the Next SHPL Installation

While the SHPL installation and demonstration aboard *USS Theodore Roosevelt* was undoubtedly a success, there are always improvements that can be made on the next installation. The following were noted in the SHPL Assessment Report:

- The Ethernet-to-ATM edge devices should be dual-homed to two ATM switches. At the time that the Cisco Catalyst 5000 units were procured, Cisco did not have dual-homed ATM cards.
- The next installation should leave more physical room for growth in each of the racks. Some of the SHPL racks on *USS Theodore Roosevelt* are completely full.
- The SHPL Network Management Workstation should be placed in a location that is most easily accessible to the network administration personnel. On *USS Theodore Roosevelt*, the NMW is in the Radio Room which is a secure space not normally accessed by SHPL maintenance personnel.
- SHPL racks should be located more central to the location of users. One of the SHPL racks on *USS Theodore Roosevelt* is so far aft that all of its fiber cable runs are quite long.
- The fiber cable backbone should be designed and installed in a mesh configuration to support the mesh-like structure of ATM. The backbone for *USS Theodore Roosevelt* was designed as a large loop throughout the ship (with FDDI in mind) and is not optimally survivable from an ATM standpoint.

7 A Vision of the Future of Shipboard Networking

7.1 Handling Voice and Data in the same Network

Many articles have been written on the issue of using ATM networks as the single integrated vehicle for handling all manner of communications – voice, video and data. But the wisdom of doing that aboard ships is still an open question, for several reasons.

First, the ATM community is in a state of transition - and uncertainty - regarding the best method of handling voice. Initially, ATM switches handled voice signals by emulating a synchronous voice circuit. This Circuit Emulation Services (CES) approach uses ATM Adaptation Layer 1 (AAL-1) and Constant Bit Rate (CBR) services which are not particularly bandwidth efficient. Each emulated voice circuit sets up a 2-way 64kbps channel and maintains that level of reserved bandwidth for the duration of the call. No other traffic can use

the bandwidth, even though one party is silent while the other party is speaking, and there are silent periods between words.

However, several ATM manufacturers have begun to offer a more bandwidth efficient technique for handling voice signals over Variable Bit Rate (VBR) services. When VBR service is combined with speech compression and silence detection, cells are transmitted through the ATM network only when needed, and data or video cells can be interleaved between the voice cells. This approach results in better ATM bandwidth utilization, and enables both voice and data to share the same virtual path/circuit (VPI/VCI). However, there are two downsides to this approach:

First, the VBR voice techniques that are currently available are proprietary. The ATM Forum is currently working on a VTOA (Voice Telephony Over ATM) standard, but it will be possibly a year to two years before products meeting the standard will be available.

Second, the VBR service does not guarantee bandwidth for the voice circuit. The cell-to-cell jitter that is typical of VBR can be compensated for by buffering, provided that the jitter is small. If the jitter becomes too large, cells must be discarded or the buffer size must be increased to the point where an unnatural delay occurs between the two sides of the conversation. In effect, with long delays, the conversation becomes like a radio circuit in which each speaker needs to say "Over" to indicate the other person's turn to speak. In a tightly-controlled ATM environment, careful load planning could assure that the jitter on the VBR voice circuits did not become too large. But if ATM is to fulfill its goal of decoupling the user from the network, and enabling multiple users to send voice, video, large data files, etc. at will, should the network be tightly managed so that voice circuits do not see excessive jitter? And if the network is not tightly managed, can we stand to have tactical voice circuits degrade in quality during critical operations which cause the network load to increase?

There is a third option for handling voice over ATM: voice over IP. Van Jacobson of the Lawrence Berkeley Laboratory at UC Berkeley developed the *Visual Audio Tool* (VAT) for supporting voice conferences over the Internet. With this technique, voice signals are digitized, optionally compressed, and then packaged into IP frames for transfer over a LAN or (more typically) a WAN. This is the technique that is used for workstation-to-workstation audio conferences on *USS Theodore Roosevelt*. After a significant effort in fine-tuning the control parameters at each end of the link, an acceptable level of voice quality using HP TAC-4 workstations with Parallax Graphics audio/video boards and Plantronics noise-canceling mic-headsets was achieved. However, this should be viewed as a convenient voice tool for workstation operators to use for talking to other workstation operators, not for general-purpose shipboard voice support. (By the way, this voice-over-IP is the same technology that is coming under fire from telephone companies. They have recently petitioned the FCC to regulate ISPs (Internet Service Providers) as telephony carriers since people are beginning to use the IP network for voice connections. When one considers the difference in quality and ease of circuit set-up between computer-based voice-over-IP and conventional POTS or ISDN voice circuits, the TELCOs really have no need to worry that ISPs are going to put them out of business.)

When you examine the difference between the shipboard networking problem and the cross-country voice-data transfer problem, it is not clear if shipboard voice should be forced onto an ATM network when it is handled so well via ISDN. ISDN handles voice with exceptional clarity, low transport delay, and rapid circuit setup time. What is the advantage in converting a synchronous sampled voice data stream into 53-byte asynchronous cells so that it can be transmitted through one or two ATM switches that span a distance of a few feet, or 1000 feet at most, to another location aboard ship where it has to be converted back to a synchronous data stream? Consider the fact that the Navy owns the fiber optic cable plant,

and does not lease it from a TELCO. Perhaps it makes more sense to use some of the fibers for ATM cell transmission and some fibers for synchronous voice transmission using ISDN/SONET.

Note that the driving function that causes industry to integrate voice and data is the extremely high cost of long-distance leased circuits. If a company with New York and San Francisco offices can give up their $10,000 per month T1 data circuit and their $20,000 per month long-distance telephone bill, and replace both with a pay-as-you-use-it integrated voice-and-data ATM access, they stand to save money. But that cost model is not true aboard ship when the ship owns the cable plant. In general, it would appear preferable to use separate fibers for voice (via ISDN/BISDN) and for data (via ATM), instead of chopping isochronous voice signals into asynchronous cells, buffering the cells to try to remove the asynchronous jitter, and then recombining cells into a quasi-isochronous data stream for voice at the other end of the circuit, when the "circuit" is at most the length of a ship. Note that this argument applies to *intra-ship* distribution of voice and data. At the interface to satellite facilities, HF facilities, dockside landlines, or other off-ship circuits, an entirely different cost/efficiency model exists and the integration of voice and data probably does make sense.

7.2 Using ATM with Traditional (Legacy) LANs

Under ideal circumstances, the Navy could make the legacy LAN to ATM transition by moving all users to ATM interfaces, perhaps 25.6 Mbps (desktop ATM) for most users and 155 Mbps for high-powered workstations. In this ideal environment, all users would have programs that wrote directly to ATM cells, all would have the benefits of quality-of-service (QoS) and bandwidth reservation, and the network would not have to emulate a LAN or deal with the problem of handling frames or frame-based addressing. However, in the real world, the user community has heavily invested in applications that are based on LAN protocols such as IP, IPX or NetBEUI, and cannot quickly change the code to work directly with ATM network interface drivers. Therefore LAN-connected users and ATM-connected users must co-exist – and communicate with each other.

There are essentially three options for handling LANs within an ATM network, and each has its own set of advantages and disadvantages.

Option 1: Multiprotocol Encapsulation using RFC-1483. RFC-1483 describes several methods of frame encapsulation, only one (LLC encapsulation) of which is popular. LLC encapsulation uses the Link Layer Control header to enable multiple protocols to share a single private virtual circuit (PVC). The RFC-1483 protocol is something that bridges and routers with ATM uplinks need to support. The ATM network itself does not need to know anything about this RFC. In effect, LLC encapsulation creates a bridge-to-bridge or router-to-router tunnel through the ATM network. While useful for tunneling through a public ATM network to create an extended private ATM network, it is of limited use aboard ship because it does not support switched virtual circuits (SVCs).

Option 2: Classical IP and ARP using RFC-1577. RFC-1577 uses the LLC encapsulation methods of RFC-1483, but adds the capability for directly-connected ATM workstations to dynamically establish connections (using SVCs) with other ATM workstations that are within the same IP subnetwork. An ATMARP (ATM address resolution protocol) server responds to ARP requests and provides ATM addresses. Workstations that use RFC-1577 hold these LAN-to-ATM address mappings in cache memory for subsequent use. Workstation that are in different IP networks cannot directly connect using RFC-1577; they need the services of a router. If an ATM router is used, it needs only a single connection to the ATM network, since it is routing between logical IP domains and not between different

network ports. Such routers are sometimes called *one-armed routers*. Note that RFC-1577 supports only IP and not any other LAN protocol.

Like RFC-1483, RFC-1577 is more directly associated with the ATM network interface cards in the workstations than with the ATM network itself. As long as an ATM network supports both PVCs and SVCs, it will support RFC-1577 communication. Of course, an ATMARP server must be provided, and this can be implemented in a workstation or within an ATM switch. If the ATM workstations are in multiple IP domains, then an ATMARP server is required for each different IP subnet.

Option 3: LAN emulation using LANE 1.0. Recognizing the limitations of RFC-1483 and RFC-1577, the ATM Forum published in May 1995 the initial version of the LAN emulation standard – LANE 1.0. The purpose of LANE is to provide transparent integration of all LAN protocols with ATM. LANE takes full advantage of switched virtual circuits (SVCs) and enables LANs and ATM-attached workstations to dynamically join emulated LANs (ELANs).

LAN emulation requires four distinct functions, some of which may be combined into the same physical unit:

> **LAN Emulation Client (LEC).** The LEC needs to be implemented in each ATM workstation that wants to join an ELAN, and in each *edge device* that is used to attach legacy LAN workstations to the ATM network. A single workstation can run more than one LEC at the same time so as to participate in multiple ELANs using the same physical interface.
>
> **LAN Emulation Server (LES).** An LES is needed for each ELAN, and it can be implemented in the *edge device*, the ATM switches, or in a separate workstation. The function of the LES is to register each new workstation into its proper ELAN, and to translate LAN MAC addresses into ATM addresses.
>
> **LAN Emulation Configuration Server (LECS).** The LECS is the overall configuration manager, and helps workstations determine which LES to talk to in a multiple-LES environment. Like the LES, the LECS can be implemented in the *edge device*, the ATM switches, or in a separate workstation. Only one LECS is needed (or can be used) in any ATM network.
>
> **Broadcast and Unknown Server (BUS).** One BUS is needed for each ELAN, and it can be implemented in any of the three devices identified above for the LES and LECS. The BUS handles the broadcast and multicast functions that are needed by legacy LANs, and also handles the initial connection of unicast transfers to workstations that are temporarily "unknown" because they are hidden behind *edge devices* and have not yet had their MAC addresses cached anywhere.

The LANE 1.0 standard defines two different frame encoding methods, one for Ethernet and one for Token Ring. (Note that both are different from the LLC encapsulation technique specified in RFC-1483.) While LANE can be used with FDDI, including FDDI's longer frame sizes, the FDDI frames must be bridged to an Ethernet or Token Ring frame structure before being handed off to LANE.

LANE has a number of advantages over RFC-1483 or RFC-1577 access to ATM. Relative to RFC-1483, its use of SVCs means that the network administrator does not have to set up all of the PVCs that would be needed to interconnect all of the station pairs if RFC-

1483 were used. Relative to RFC-1577, LANE has the advantage of operating at the link layer, not the network layer, and therefore working equally well with IP, IPX, DECnet, and even non-routeable protocols such as LAT or NetBEUI. Furthermore, RFC-1577 works only with directly-attached ATM devices; it cannot interconnect ATM workstations with LAN workstations without the services of a router that has both LAN and ATM interfaces.

In spite of these advantages, LANE 1.0 is not the final solution to the ATM integration of legacy LANs. The LANE 1.0 standard does not define how multiple LESs and multiple BUSs can be used to provide redundancy. The system administrator is faced with four approaches for adding redundancy to an ATM-based virtual LAN environment: (1) Using proprietary LAN emulation protocols exclusively; (2) using standards-based LANE components which are single-point failures on the network; (3) using standards-based LANE components that have proprietary extensions for supporting redundancy (e.g., Fore Systems network interface cards that were used in SHPL); or (4) implementing redundant LANE components in a non-standard but survivable way (as was done for SHPL). An update called LANE 2.0 is in the works and it is supposed to address some of these redundancy concerns.

Another limitation of LANE is that of broadcast performance. Even though all new workstation network connections, even for unicast, require a time-consuming processing by the LECS and LES, once this is done the needed ATM addresses are cached and subsequent connections can be made rapidly. For broadcast and multicast transfers, however, the BUS will always be used, and its performance (throughput and delay) directly impact the network. In a network with a high broadcast/multicast data load, the limitation of one BUS per ELAN may create serious bottlenecks.

Another ATM Forum effort that is closely related to LANE is MPOA – multiprotocol over ATM. MPOA will operate at the network layer in addition to the link layer, and will be able to route multiple protocols, not just IP. MPOA is still one or two years away from completion, however.

7.3 Frame versus Cell Based Applications

Most legacy LAN applications used aboard Navy ships today are written based upon the assumption that the network handles *frames*, e.g., Ethernet, Token Ring, FDDI, etc. and that the transport mechanism is either IP or IPX. The optimum solution for ATM networks is for applications to write to *cells* and to exploit all of the low-delay, QoS (quality of service) and bandwidth advantages of cell-based transfer. In the CVN-71 network, the only application that uses cell-based transfer is the Real Time Display Software (RTDS) used for the distribution of TV signals via the ATM network. All other applications use IP, primarily because they were developed during an era in which IP was the *de facto* protocol of choice. The Navy needs to invest in the development of new applications that exploit the full capabilities of cell-based interface to ATM, and also to port existing applications from frame-based protocols to cell-based. Not only will direct cell-based applications be able to exploit QoS and bandwidth reservation, but the added delay of frame-to-cell and cell-to-frame conversion can be avoided.

Porting existing applications to work with cell-based network drivers, or writing new applications to an ATM application programming interface, does not necessarily mean that all workstations will have to interface at the 155 Mbps OC-3c level. The 25.6 Mbps ATM-25 or Desktop ATM interface offers a lower-cost option for workstations that do not need, or cannot effectively use, the 155 Mbps bandwidth. An even lower cost option that promises all of the ATM advantages of bandwidth reservation and quality of service is the *cells in frames* technique that uses an ATM cell structure riding on top of the standard Ethernet media access

control (MAC) layer. With *cells in frames* (CIF), workstations that are already equipped with Ethernet cards can continue to use these network cards and switch to ATM cell-based applications through the use of a CIF driver between the application and the Ethernet card. Of course, CIF-ATM edge devices will be needed to convert the Ethernet/CIF interfaces to true ATM interfaces for interface to ATM backbone switches.

7.4 A Preferred Follow-on SHPL Architecture

As noted previously, while the SHPL installation on *USS Theodore Roosevelt* was definitely a success, there are improvements that could be made on the next shipboard installation. A preferred architecture for a follow-on shipboard installation of SHPL is shown in Figure 4.

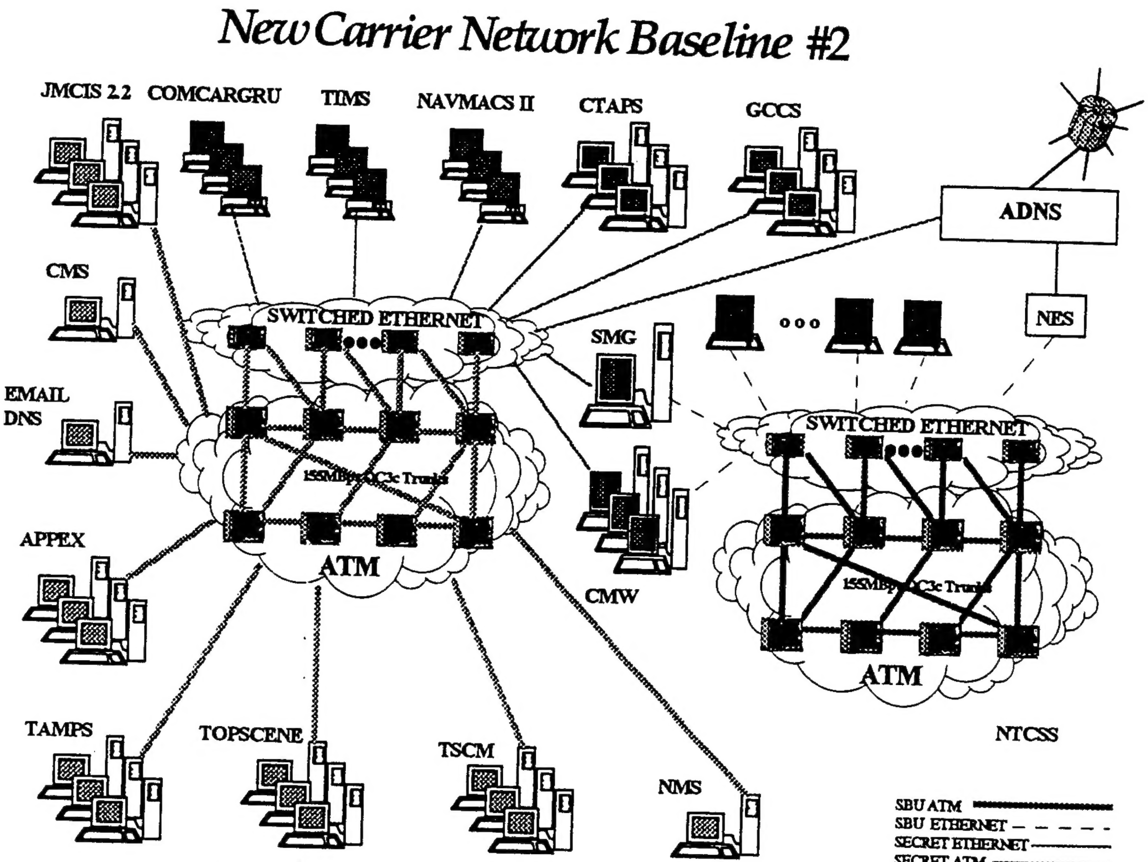

Figure 4. A Proposed Follow-on Architecture for SHPL

The significant differences between this follow-on diagram and the current architecture shown in Figure 1 are as follows: (a) JMCIS moves from an Ethernet/WhisperLan connection to ATM, (b) the Cable Management System (CMS) and email/DNS machines move from Ethernet to ATM, and (c) the NTCSS system moves from an FDDI backbone with shared-Ethernet workstations to an ATM backbone with switched-Ethernet workstations.

7.5 The Long-Range Architecture

The general direction that we see shipboard networking moving toward is suggested by the diagram of Figure 5. Here the changes are more significant: (a) There is only one

ATM backbone; it is multilevel secure (MLS) and serves all shipboard users, at least up to the GENSER SECRET level. (b) The voice system is integrated into the ATM backbone via ISDN switches which are interconnected via ATM. The voice signals may be carried via Circuit Emulation Services using AAL-1 and Constant Bit Rate, or possibly via voice-over-VBR services using AAL-5. (c) There are three levels of ATM-aware workstations: high-performance workstations that interface at the 155 Mbps OC-3c level, medium performance workstations that interface at the 25.6 Mbps "Desktop 25" level, and low cost workstations that interface via 10 Mbps fiber Ethernet. The latter fall into two groups: ordinary Ethernet workstations that run applications over IP, and "cells in frames" workstations that run ATM-aware applications over the Ethernet MAC layer but with an ATM driver using the *cells in frames* technology.

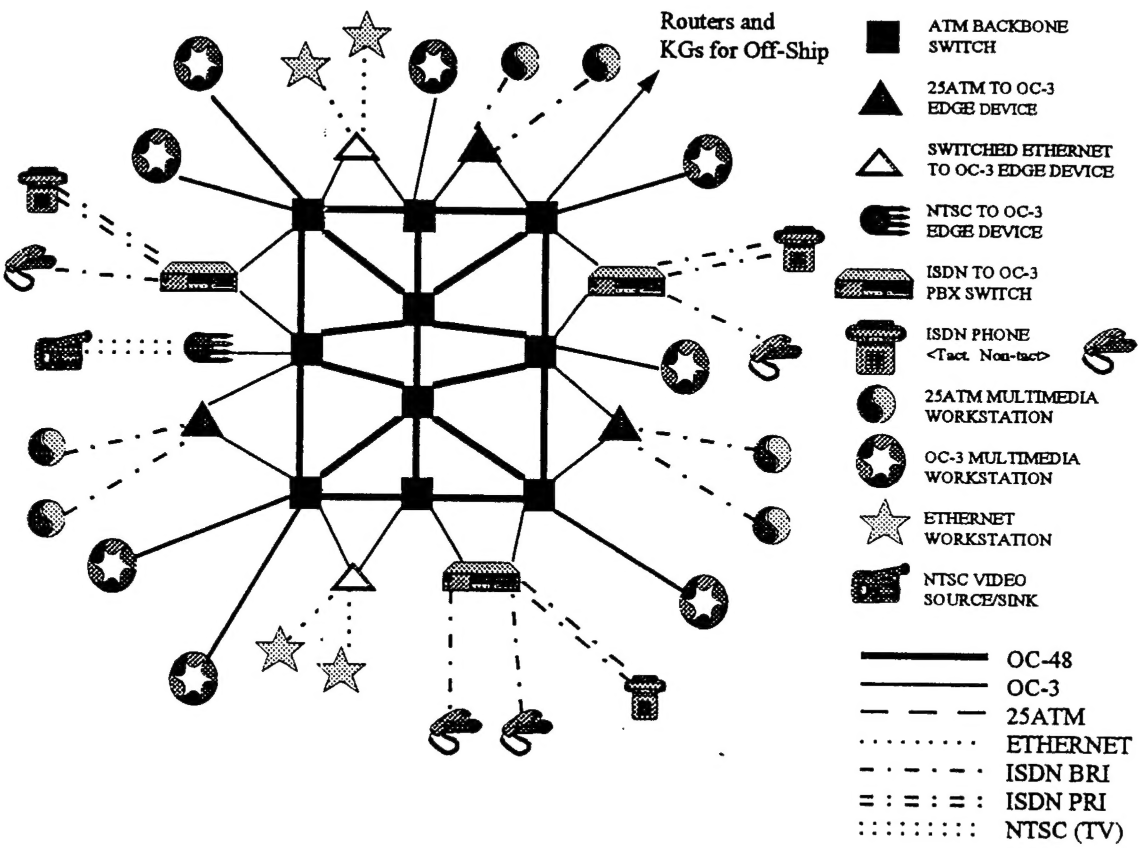

Figure 5. A Long-Range View of Shipboard Networking

Whether shipboard networking actually moves to this single-backbone architecture for data, video and voice, or maintains a separate ISDN-SONET voice system that does not transit the ATM switches, is still an open question that requires further study at both the technology level and the logistics support level. In either case, the ATM network installed on *USS Theodore Roosevelt* will be remembered as a seminal effort that established beyond doubt that commercial ATM technology can survive, and can be effectively used for high-bandwidth tactical applications, aboard Navy ships.